

IN THE CLAIMS:

Please amend the claims as follows. The claims are in the format as required by 35 C.F.R. § 1.121.

1. (Currently amended) A method for regulating resource usage by a plurality of distributed applications running on a plurality of interconnected machines, the method comprising:

receiving an initial resource policy having a set of resource utilization goals;

gathering, by a plurality of interconnected machines, information on the state and health of applications and system resources at each of the machines;

providing each machine receiving a subset of the total information gathered and a portion of [[a]] the initial global, distributed resource policy specifying allocation of resources of the plurality of machines to the set of distributed applications, wherein the subset of the total information received by each machine corresponds to one or more applications running on that machine;

determining resources available at the plurality of interconnected machines;

each machine detecting one or more requests for resources by each of the plurality of distributed applications running on the plurality of interconnected machines;

at each of the plurality of interconnected machines, allocating the resources to each of the plurality of distributed applications based on the subset of the total information gathered;

determining the resource consumption by the plurality of each application based on its actual usage;

periodically exchanging resource information amongst the plurality of interconnected machines, wherein the resource information includes updated resource policies, requests for the resources and resource availability at at least one of the plurality of interconnected machines; and

at each of the plurality of interconnected machines, allocating adjusting the resources to each of the plurality of distributed applications based upon the global, distributed resource policy and the resource information periodically exchanged resource information received by that machine.

2. (Currently amended) The method of claim 1, wherein said resources include one or more of communication resources, network bandwidth shared by the plurality of machines, computing resources, processing resources available at the plurality of machines, memory

resources available at the plurality of machines.

3 – 12. (Cancelled).

13. (Currently amended) The method of claim ~~[[12]]~~ 1, wherein each machine comprises:  
a kernel space; and

a user space module, wherein each of the user-space modules comprises:

a data store for writing persistent data to the local disc of the host  
computer, wherein data is redundantly replicated on the local disc of a  
configurable number of peer hosts;

a shepherd for organizing and maintaining a configuration of machines in  
the kernel space which can be made available to other resources;

a policy manager for tracking the policies specified for applications and  
the available resources, for monitoring the resource utilization, and for storing  
the monitoring information in the data store;

a request manager for authenticating a user of the system and for  
passing valid requests received from the user to the user-space module;

an event scheduler for receiving events from the kernel space and  
supplying the information to other components in the user-space module; and

a process scheduler for receiving information from the policy manager about resource  
entitled to each program, monitoring all processes running on a machine and communicating to  
the policy manager the resource utilization of each program,

wherein the priorities for allocation of resources are automatically adjusted based on  
occurrence of particular events, wherein the policy manager receives information about the  
occurrence of the events and monitors the resource utilization of all applications, jobs,  
processes and flows and records the information in a data store.

14. (Original) The method of claim 1, wherein said detecting step includes detecting each instance of a program running at each of the plurality of machines.

15. (Currently amended) The method of claim 1, further comprising a system user  
interface, wherein said detecting step is performed by application detection rules with a  
frequency established by a user.

16. (Cancelled).

17. (Currently amended) The method of claim 1, wherein the plurality of interconnected machines comprises a flow scheduler, wherein each machine includes a portion of the flow scheduler, wherein said exchanging step is performed by the flow scheduler with a frequency established by a user.

18-20. (Cancelled).

21. (Original) The method of claim 1, wherein said exchanging step includes using a bandwidth-conserving protocol.

22. (Currently amended) The method of claim 1, wherein said allocating step includes one or more of regulating usage of resources by each of the plurality of programs, scheduling processing resources at each of the plurality of machines, regulating the volume of communications sent by a particular program and delaying the sending of a communication by a particular program.

23-25. (Cancelled).

26. (Original) The method of claim 1, further comprising:  
collecting resource information regarding requests for resources and resource availability;  
and  
generating resource utilization information for display to a user based upon the collected resource information.

27-29. (Cancelled).

30. (Currently amended) A system for regulating utilization of computer resources of a plurality of interconnected computers, the system comprising:  
a plurality of computers having local resources to be regulated which are interconnected to each other through a network, wherein each computer comprises:  
a kernel space; and

a user-space module comprising:

a monitoring module provided at each computer having resources to be regulated, for monitoring resource utilization and providing the resource utilization information to other interconnected computers having resources to be regulated;

a manager module for providing a global, distributed policy governing the global utilization of resources available on the plurality of computers; and

an enforcement module at each computer for which resources are to be regulated for regulating usage of resources based on said global, distributed policy and the resource utilization information received from other interconnected computers,

wherein each computer of the plurality of computers periodically exchanges information with other computers of the plurality of computers, including updates to the global, distributed policy.

31-32. (Cancelled).

33. (Currently amended) The system of claim 30, wherein said ~~monitoring module at a given computer identifies at least one application running at the given computer~~ user-space module comprises:

a data store for writing persistent data to the local disc of the host computer, wherein data is redundantly replicated on the local disc of a configurable number of peer hosts;

a shepherd for organizing and maintaining a configuration of machines in the kernel which can be made available to other resources;

a policy manager for tracking the policies specified for applications and the available resources, for monitoring the resource utilization, and for storing the monitoring information in the data store;

a request manager for authenticating a user of the system and for passing valid requests received from the user to the user-space module;

an event scheduler for receiving events from the kernel space and supplying the information to other components in the user-space module; and

a process scheduler for receiving information from the policy manager about resource entitled to each program, monitoring all processes running on a machine and communicating to the policy manager the resource utilization of each program.

34-39. (Cancelled).

40. (Currently amended) The system of claim 30, wherein said monitoring module at a given computer provides resource utilization information to each other connected computer in response to particular events, wherein said resource utilization information provided by said monitoring module includes one or more of information regarding requests for communication resources and information regarding requests for processing resources.

41. (Currently amended) The system of claim 30, wherein said resource utilization information provided by said monitoring module includes information regarding requests for one or more of communication resources and processing resources.

42-45. (Cancelled).

46. (Original) The system of claim 30, wherein said monitoring module uses a bandwidth-conserving protocol for providing resource utilization information.

47. (Original) The system of claim 30, wherein said enforcement module schedules processing resources at each of the plurality of computers based on said transferred rules and the resource utilization information.

48-49. (Cancelled).

50. (Currently amended) The system of claim 30, further comprising:  
a configuration module for a user to establish rules governing utilization of resources, wherein said configuration module collects resource utilization information from the plurality of computers, wherein said configuration module suggests rules governing utilization of resources based, at least in part, upon the collected resource utilization information, and wherein said

configuration module is operable to display the collected resource utilization information to a user.

51-53. (Cancelled).

54. (Currently amended) A method for scheduling communications by a plurality of applications running on a plurality of computers connected to each other through a network, the method comprising:

gathering, by a plurality of interconnected machines, information on the state and health of applications and system resources at each of the machines;

providing receiving, at each of the interconnected machines, a subset of a global, distributed policy, wherein the policy specifying specifies priorities for scheduling communications by the plurality of applications;

~~periodically determining communication resources available for sharing at the plurality of computers;~~

at each of the plurality of computers, detecting requests to communicate and identifying a particular application associated with each request;

at each of the plurality of interconnected machines, allocating the resources to each of the plurality of distributed applications based on the subset of the total information gathered;

determining the resource consumption by the plurality of each application based on its actual usage;

periodically exchanging bandwidth information amongst the plurality of computers, wherein the bandwidth information includes the application making the requests to communicate and a measure of the communications resources required to fulfill the requests;  
and

at each of the plurality of computers, scheduling communications based upon the received subset of the global, distributed policy and the bandwidth information.

55. (Cancelled).

56. (Original) The method of claim 54, wherein said communication resources include network bandwidth shared by the plurality of computers.

57-58. (Cancelled).

59. (Currently amended) The method of claim 54, wherein said scheduling step includes one or more of the steps of immediately transmitting all communications if the bandwidth information indicates communication traffic is light, delaying a portion of the communications if the bandwidth information indicates communication traffic is heavy, delaying transmission of communications by lower-priority applications and load balancing, wherein said load balancing includes redirecting communications received at a first computer to a second computer.

60-65. (Cancelled).